arXiv:1604.00912v1 [stat.ME] 4 Apr 2016

A multivariate nonlinear mixed effects model for longitudinal image analysis: Application to amyloid imaging[☆,☆☆]

Murat Bilgel[a,b,c,*], Jerry L. Prince[a,b,d,e], Dean F. Wong[e], Susan M. Resnick[c], Bruno M. Jedynak[f]

[a] *Image Analysis and Communications Laboratory, Johns Hopkins University School of Engineering, Baltimore, MD, USA*
[b] *Dept. of Biomedical Engineering, Johns Hopkins University School of Medicine, Baltimore, MD, USA*
[c] *Laboratory of Behavioral Neuroscience, National Institute on Aging, National Institutes of Health, Baltimore, MD, USA*
[d] *Dept. of Electrical and Computer Engineering, Johns Hopkins University School of Engineering, Baltimore, MD, USA*
[e] *Dept. of Radiology, Johns Hopkins University School of Medicine, Baltimore, MD, USA*
[f] *Dept. of Mathematics and Statistics, Portland State University, Portland, OR, USA*

Abstract

It is important to characterize the temporal trajectories of disease-related biomarkers in order to monitor progression and identify potential points of intervention. These are especially important for neurodegenerative diseases, as therapeutic intervention is most likely to be effective in the preclinical disease stages prior to significant neuronal damage. Neuroimaging allows for the measurement of structural, functional, and metabolic integrity of the brain at the level of voxels, whose volumes are on the order of mm^3. These voxelwise measurements provide a rich collection of disease indicators. Longitudinal neuroimaging studies enable the analysis of changes in these voxelwise measures. However, commonly used longitudinal analysis approaches, such as linear mixed effects models, do not account for the fact that individuals enter a study at various disease stages

[☆]DOI: 10.1016/j.neuroimage.2016.04.001
[☆☆]This manuscript version is made available under the CC-BY-NC-ND 4.0 license `http://creativecommons.org/licenses/by-nc-nd/4.0/`.
*Corresponding author at: National Institute on Aging, Laboratory of Behavioral Neuroscience, 251 Bayview Blvd., Suite 100, Rm 04B316, Baltimore, MD 21224, USA. Phone: +1-410-558-8151. Fax: +1-410-558-8674.
Email address: `murat.bilgel@nih.gov` (Murat Bilgel)

and progress at different rates, and generally consider each voxelwise measure independently. We propose a multivariate nonlinear mixed effects model for estimating the trajectories of voxelwise neuroimaging biomarkers from longitudinal data that accounts for such differences across individuals. The method involves the prediction of a progression score for each visit based on a collective analysis of voxelwise biomarker data within an expectation-maximization framework that efficiently handles large amounts of measurements and variable number of visits per individual, and accounts for spatial correlations among voxels. This score allows individuals with similar progressions to be aligned and analyzed together, which enables the construction of a trajectory of brain changes as a function of an underlying progression or disease stage. We apply our method to studying cortical β-amyloid deposition, a hallmark of preclinical Alzheimer's disease, as measured using positron emission tomography. Results on 104 individuals with a total of 300 visits suggest that precuneus is the earliest cortical region to accumulate amyloid, closely followed by the cingulate and frontal cortices, then by the lateral parietal cortex. The extracted progression scores reveal a pattern similar to mean cortical distribution volume ratio (DVR), an index of global brain amyloid levels. The proposed method can be applied to other types of longitudinal imaging data, including metabolism, blood flow, tau, and structural imaging-derived measures, to extract individualized summary scores indicating disease progression and to provide voxelwise trajectories that can be compared between brain regions.

Keywords: Longitudinal image analysis, progression score, amyloid imaging

1. Introduction

It is important to characterize the temporal trajectories of disease-related biomarkers in order to monitor progression and to identify potential points of intervention. Such a characterization is especially important for neurodegenerative diseases, as therapeutic intervention is most likely to be effective in the preclinical disease stages prior to significant neuronal damage. For example, in Alzheimer's

disease, brain changes evident in structural, functional, and metabolic imaging may occur more than a decade before the onset of cognitive symptoms (Bateman et al., 2012), with cortical amyloid-β (Aβ) accumulation being one of the earliest changes (Jack et al., 2013; Sperling et al., 2014a; Villemagne et al., 2013). Such brain changes can be measured using neuroimaging techniques and can be tracked over time at the individual level via longitudinal studies.

Given the focus on preventing and delaying the onset of incurable neurodegenerative diseases, the emphasis of clinical trials has shifted to studying clinically normal individuals with positive biomarkers, for example those exhibiting brain amyloid in the case of AD, in order to identify early intervention opportunities in the preclinical stages of disease (Sperling et al., 2014b). It is important to determine the temporal trajectories of hypothesized biomarkers in the early disease stages in order to better understand their associations with disease progression. Current neuroimaging methods allow for the characterization of the brain at the mm^3 level, generating hundreds of thousands of measurements that can be used as potential biomarkers of neurodegenerative diseases. Understanding the temporal trajectories of these voxelwise measurements can provide clues into disease mechanisms by identifying the earliest and fastest changing brain regions.

Changes in voxelwise neuroimaging measurements over time are commonly studied using linear mixed effects models (Bernal-Rusiel et al., 2012, 2013; Ziegler et al., 2015). Univariate linear mixed effects models use time or age to characterize changes in a single imaging measure. However, time or age may not be the appropriate metric for measuring disease progression due to variability across individuals. While covariates can be included in linear mixed effects models to account for this variability, choosing the correct set of covariates is difficult and covariates generally have a more complicated association with disease progression than the assumed linear relationship of linear mixed effects models. Instead, this variability can be accounted for by aligning individuals in time based on their longitudinal biomarker profiles within a multivariate framework. This is the premise of the Disease Progression Score method, which has been applied to studying changes in cognitive and biological markers related

to Alzheimer's disease (Jedynak et al., 2012, 2014; Bilgel et al., 2014). It is assumed that there is an underlying progression score (PS) for each subject visit that is an affine transform of the subject's age, and given this PS, it is possible to place biomarker measurements across a group of subjects onto a common timeline. The affine transformation of age removes across-subject variability in baseline biomarker measures as well as in their rates of longitudinal progression. Each biomarker is associated with a parametric trajectory as a function of PS, whose parameters are estimated along with the PS for each subject. This allows one to "stitch" data across subjects to obtain temporal biomarker trajectories that fit an underlying model (Fig. 1).

Previous approaches have used certain cognitive measures, such as ADAS-Cog (Caroli and Frisoni, 2010; Yang et al., 2011), MMSE (Doody et al., 2010) or CDR-SB (Delor et al., 2013) as a surrogate for disease progression to delineate the trajectories of other AD-related cognitive measurements. These methods operate with the assumption that disease progression is reflected by a single cognitive measurement rather than a profile of multiple measurements, and therefore are inherently limited in their characterization of disease evolution. Younes et al. (2014) fitted a piecewise linear model to longitudinal data assuming that each biomarker becomes abnormal a certain number of years before clinical diagnosis, and this duration was estimated for each biomarker to yield longitudinal trajectories as a function of time to diagnosis. A quantile regression approach was employed by Schmidt-Richberg et al. (2015) to align a sample of cognitively normals and mild cognitively impaired (MCI) with a sample of MCI and AD, and then to estimate biomarker trajectories. These approaches assume that all individuals are on a path to disease and require knowledge of clinical diagnosis. Therefore, they are not suitable for studying the earliest changes in individuals who have not converted to a clinical diagnosis. Donohue et al. (2014) applied a self-modeling regression model within a multivariate framework to characterize the longitudinal trajectories of a set of cognitive, CSF, and neuroimaging-based biomarkers. This approach allows for across-subject variability only in the age of onset, not in progression speed. Models incorporating fixed effects as well

as individual-level random effects have been proposed to study ADAS-Cog (Ito et al., 2011; Schiratti et al., 2015b) and regional cortical atrophy (Schiratti et al., 2015b), and Schulam et al. (2015) used a spline model that incorporates longitudinal clustering and modeling of individual-level effects to study trajectories of scleroderma markers. These mixed effects models take into consideration each measure separately rather than using them within a unifying framework. Others have used event-based probabilistic frameworks to determine the ordering of changes in longitudinal biomarker measures as well as the appropriate thresholds for separating normal from abnormal measures (Fonteijn et al., 2012; Young et al., 2014). These methods characterize longitudinal biomarker trajectories in a discrete framework rather than a continuous one. Schiratti et al. (2015a) proposed an extension to their earlier approach to model multiple measures together. Biomarker trajectories are assumed to be identical except for a shift along the disease timeline, and this assumption prevents hypothesis testing regarding rate of change across biomarkers. Furthermore, biomarkers are assumed to be conditionally independent given the subject-level random effects, but this assumption is not realistic when biomarkers are voxel-based neuroimaging measurements.

Here, we adapt the disease progression score principle to studying longitudinal neuroimaging data by making substantial innovations to the progression score model and parameter estimation procedure. First, voxelwise imaging measures constitute the biomarkers in the model, and are analyzed together in a multivariate framework. Studying progression at the voxel level rather than using region of interest (ROI)-based measures allows for the discovery of patterns that may not be confined within any given ROI. Second, since voxelwise imaging measures have an underlying spatial correlation, we incorporate the modeling of the spatial correlations among the biomarker error terms. Modeling spatial correlations makes the inference of the subject-specific progression scores less susceptible to the inherent correlations among the voxels. Third, we incorporate a bivariate normal prior on the subject-specific variables that define the relationship between age and PS. The prior allows a better modeling of the variance

within individuals and enables the incorporation of individuals with a single visit into the model fitting procedure. Fourth, instead of using an alternating least-squares approach for parameter estimation as presented by Jedynak et al. (2012), we formulate the model fitting as an expectation-maximization (EM) algorithm, which guarantees convergence to a local maximum and allows for an efficient model fitting framework for a large number of biomarkers. Finally, we present a statistical framework for comparing the onset and rate of progression across different regions. This paper extends our previous approach for analyzing longitudinal voxelwise imaging measures using the progression score framework by incorporating a prior on the subject-specific variables, presenting a hypothesis testing framework for determining biomarker ordering, and performing an extensive validation of the method (Bilgel et al., 2015c).

We first show using simulated data that the model parameters are estimated accurately and that modeling spatial correlations improves parameter estimation. We then apply the method to distribution volume ratio (DVR) images derived from Pittsburgh compound B (PiB) PET imaging, which show the distribution of cerebral fibrillar amyloid. Models fitted using data for 104 participants with a total of 300 PiB-PET visits reveal that the precuneus and frontal cortex show the greatest longitudinal increases in fibrillar amyloid, with smaller increases in lateral temporal and temporoparietal regions, and minimal increases in the occipital cortex and the sensorimotor strip. Our results suggest that the precuneus is the earliest cortical region to accumulate amyloid. The results are consistent across the two hemispheres, and the estimated PS agrees with a widely used PET-based global index of brain amyloid known as mean cortical DVR. The presented method can be applied to other types of longitudinal imaging data to understand voxelwise trajectories and to quantify each individual scan against the estimated progression pattern.

2. Method

2.1. Model

Our goal is to characterize the progression of disease or an underlying process as measured using a collection of relevant biomarkers. Disease or process stage, as indicated by a progression score (PS), s, is intrinsically related to time t, measured as the age of a subject. Since individuals differ in their onset and rate of progression, the relationship between s and t varies across individuals. We model the progression s as a linear function of time t for each individual and allow for the prediction of separate slopes and intercepts to account for this variability across individuals.

Generally, there is a particular presentation of symptoms and biomarker measurements at a given progression stage. Furthermore, as the disease or process progresses, there is a particular temporal progression of the biomarkers. In this work, we consider voxelwise PET measures as biomarkers and model the temporal trajectory of each biomarker, or voxel, as a linear function of the progression s. Considering voxelwise neuroimaging measures as biomarkers may appear unusual; however, these measures fit the NIH definition of a biomarker: "a characteristic that is objectively measured and evaluated as an indicator of normal biological processes, pathogenic processes, or pharmacologic responses to a therapeutic intervention" (Biomarkers Definitions Working Group, 2001). As illustrated in Figure 1, PS aligns longitudinal biomarker measures better than age since it accounts for differences across individuals in rates as well as baseline levels of progression. After this alignment in time, the estimated biomarker trajectories can be compared on the common PS scale.

In the following subsections, we describe the progression score model in detail.

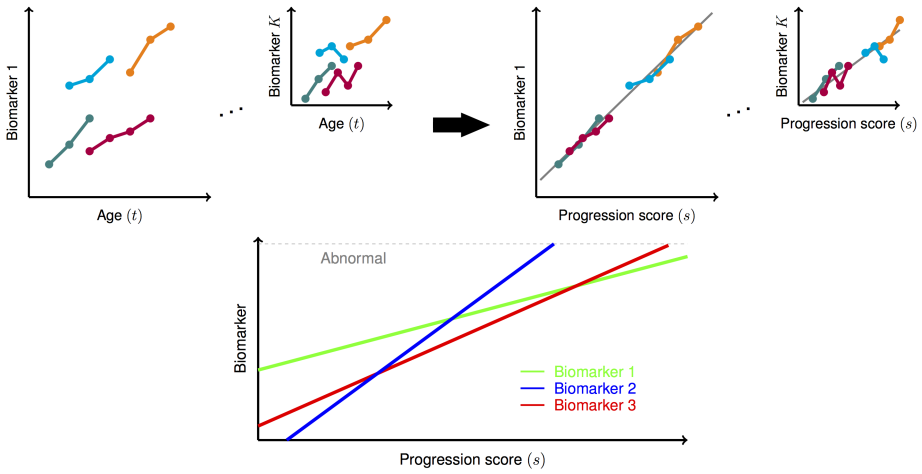

Figure 1: Illustration of the biomarker alignment concept in the progression score model. The biomarkers we consider in this work are PET measures of cerebral amyloid across a total of $K \approx 30,000$ voxels. *Top:* Progression score (PS) aligns longitudinal measures better than age, and allows for the estimation of a trajectory for each biomarker/voxel (in gray). *Bottom:* Estimated biomarker trajectories can be compared on the common PS scale.

2.1.1. Subject-specific model

The progression score s_{ij} for subject i at visit j is assumed to be an affine transformation of the subject's age t_{ij}:

$$\begin{aligned} s_{ij} &= \alpha_i t_{ij} + \beta_i \\ &= \mathbf{q}_{ij}^T \mathbf{u}_i, \end{aligned} \tag{1}$$

where $\mathbf{q}_{ij} = \begin{bmatrix} t_{ij} \\ 1 \end{bmatrix}$, and $\mathbf{u}_i = \begin{bmatrix} \alpha_i \\ \beta_i \end{bmatrix}$. The subject-specific variables, α_i and β_i contained in the vector $\mathbf{u}_i$, are assumed to follow a bivariate normal distribution, i.e., $\mathbf{u}_i \sim \mathcal{N}_2(\mathbf{m}, V)$, which are independent and identically distributed across subjects. This model accounts for differences between subjects in the rate of progression via α, and in the baseline levels of disease progression via β.

2.1.2. Subject-specific prior covariance model

The prior covariance V is modeled as a 2×2 unstructured covariance matrix. Log-Cholesky parametrization of V, given by $\boldsymbol{\nu}$, ensures that $V \equiv V(\boldsymbol{\nu})$ is positive

definite (Pinheiro and Bates, 1996). Let $U = \begin{bmatrix} U_{11} & U_{12} \\ 0 & U_{22} \end{bmatrix}$ be an upper triangular matrix such that $V = U^T U$. If the diagonal elements of U are constrained to be positive, then this Cholesky decomposition is unique. To ensure that the diagonal elements of U are positive in an unconstrained optimization framework, we use the natural logarithm of the diagonal elements of U as parameters. We then vectorize the upper triangular elements (including the diagonal) of U to obtain the parameter vector $\boldsymbol{\nu} = [\log U_{11} \;\; U_{12} \;\; \log U_{22}]^T$, which uniquely parameterizes V and ensures its positive definiteness.

2.1.3. Biomarker trajectory model

The collection of K biomarker measurements form the $K \times 1$ vector $\mathbf{y}_{ij}$ for subject i at visit j. Longitudinal trajectories associated with these biomarkers are assumed to be linear and parameterized by $K \times 1$ vectors $\mathbf{a}$ and $\mathbf{b}$:

$$\mathbf{y}_{ij} = \mathbf{a}\, s_{ij} + \mathbf{b} + \boldsymbol{\epsilon}_{ij}. \tag{2}$$

Here, $\mathbf{a} = [a_1, a_2, \ldots, a_K]^T$, $\mathbf{b} = [b_1, b_2, \ldots, b_K]^T$, and $\boldsymbol{\epsilon}_{ij} \sim \mathcal{N}_K(0, R)$ is the observation noise. $\boldsymbol{\epsilon}_{ij}$ are assumed to be independent and identically distributed across subjects and visits.

2.1.4. Noise covariance model

The matrix R is assumed to have the form $R = \Lambda C \Lambda$, where Λ is a diagonal matrix with positive diagonal elements $\boldsymbol{\lambda}$ and C is a correlation matrix parameterized by $\boldsymbol{\rho}$. This parameterization guarantees that R is a positive definite matrix (Galecki and Burzykowski, 2013). For ease of notation, we let $\Lambda \equiv \Lambda(\boldsymbol{\lambda})$, $C \equiv C(\boldsymbol{\rho})$, and $R \equiv R(\boldsymbol{\lambda}, \boldsymbol{\rho})$.

When the biomarkers under consideration have a spatial organization, i.e., if they are voxelwise measurements from medical images, then the correlation matrix C can be described as a function of the spatial distance $d \equiv d(k, k')$ between pairs of voxels indexed by k and k' as well as the spatial correlation parameter $\boldsymbol{\rho}$. Possible univariate parameterizations (i.e., $\boldsymbol{\rho} = \rho \in \mathbb{R}$) of C are

presented in Table 1. All of these spatial correlation functions ensure that C is a valid correlation matrix (Galecki and Burzykowski, 2013).

Table 1: Spatial correlation functions.

Exponential	$C_{kk'} = e^{-d/\rho}$
Gaussian	$C_{kk'} = e^{-(d/\rho)^2}$
Rational quadratic	$C_{kk'} = \frac{1}{1+(d/\rho)^2}$
Spherical	$C_{kk'} = \left(1 - \frac{3}{2}\frac{d}{\rho} + \frac{1}{2}\left(\frac{d}{\rho}\right)^3\right) \mathbb{1}(d < \rho)$

$d \equiv d(k, k')$ is the spatial distance between voxels indexed by k and k'.

2.1.5. Overall model

The overall model, diagrammatically summarized using plate notation in Fig. 2, is described by the following equations:

$$s_{ij} = \mathbf{q}_{ij}^T \mathbf{u}_i \tag{3}$$

$$\mathbf{y}_{ij} = \mathbf{a}\, s_{ij} + \mathbf{b} + \boldsymbol{\epsilon}_{ij} \tag{4}$$

$$\mathbf{u}_i \sim \mathcal{N}_2(\mathbf{m}, V(\boldsymbol{\nu})) \tag{5}$$

$$\boldsymbol{\epsilon}_{ij} \sim \mathcal{N}_K(0, R(\boldsymbol{\lambda}, \boldsymbol{\rho})). \tag{6}$$

While this model is a mixed effects model since it incorporates the fixed effects $\mathbf{a}, \mathbf{b}$ as well as the individual-level random effects $\mathbf{u}_i$, and is nonlinear in the parameters, it departs from the form of the nonlinear mixed effects model described by Lindstrom and Bates (1990). Therefore, instead of pursuing a restricted maximum likelihood approach, we use an expectation-maximization (EM) approach, as described below.

Let $\boldsymbol{\theta}$ be the collection of model parameters $\mathbf{m}, \boldsymbol{\nu}, \mathbf{a}, \mathbf{b}, \boldsymbol{\lambda}, \boldsymbol{\rho}$. The complete

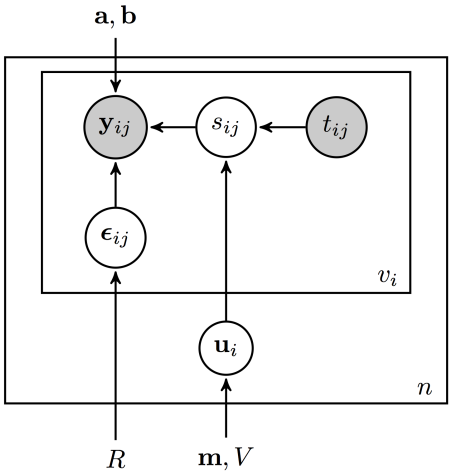

Figure 2: Probabilistic model of the progression score using plate notation. The progression score s_{ij} establishes the link between age t_{ij} and the voxelwise observations $\mathbf{y}_{ij}$. Circles indicate variables, and known measures are shaded. Rectangles indicate that the same model applies across visits (inner rectangle) and individuals (outer rectangle). Arrows indicate dependencies between variables and parameters.

log-likelihood for the model is

$$\begin{aligned}
\ell(\mathbf{y}, \mathbf{u}; \boldsymbol{\theta}) &= \sum_i \ell(\mathbf{y}_i, \mathbf{u}_i; \boldsymbol{\theta}) && (7)\\
&= \sum_i \log f(\mathbf{y}_i \mid \mathbf{u}_i; \boldsymbol{\theta}) + \log f(\mathbf{u}_i; \boldsymbol{\theta}) && (8)\\
&= -\frac{1}{2}\sum_{i,j} \log|2\pi R| - \frac{1}{2}\sum_{i,j} (\mathbf{y}_{ij} - Z_{ij}\mathbf{u}_i - \mathbf{b})^T R^{-1} (\mathbf{y}_{ij} - Z_{ij}\mathbf{u}_i - \mathbf{b}) \\
&\quad -\frac{1}{2}\sum_{i} \log|2\pi V| - \frac{1}{2}\sum_{i} (\mathbf{u}_i - \mathbf{m})^T V^{-1} (\mathbf{u}_i - \mathbf{m}), && (9)
\end{aligned}$$

where $Z_{ij} = \mathbf{a}\mathbf{q}_{ij}^T$. $\mathbf{y}_i$ is a vector of all biomarker measures stacked across all visits of individual i, and $\mathbf{y}_{ij}$ is a vector of all biomarker measures for individual i at visit j. Given n individuals, we use the summation notation $\sum_{i,j}$ as a shorthand for $\sum_{i=1}^{n}\sum_{j=1}^{v_i}$, where v_i is the number of visits for individual i.

Since the marginal likelihood $f(\mathbf{y}; \theta)$ involves an integral over all possible values of $\mathbf{u}_i$, maximizing it directly is difficult. Therefore, we use the expectation-maximization (EM) approach, where we consider the subject-specific variables

$\{\mathbf{u}_i\}$ as hidden. EM allows us to reformulate the maximization problem in terms of the complete log-likelihood $\ell(\mathbf{y}, \mathbf{u}; \theta)$. The observations $\mathbf{y}$ include biomarker measurements $\{\mathbf{y}_{ij}\}$ at each visit. The unknown parameters are the subject-specific variable distribution parameters $\mathbf{m}$ and $\boldsymbol{\nu}$, the trajectory parameters $\mathbf{a}$ and $\mathbf{b}$, and the noise covariance parameters $\boldsymbol{\lambda}$ and $\boldsymbol{\rho}$.

2.2. E-step

$(\mathbf{y}, \mathbf{u})$ are the complete data. Let $\boldsymbol{\theta}' = \{\mathbf{m}', \boldsymbol{\nu}', \mathbf{a}', \mathbf{b}', \boldsymbol{\lambda}', \boldsymbol{\rho}'\}$ be the previous parameter estimates. By Proposition A.1, the E-step integral is proportional to $\sum_i \int \Phi(\tilde{\mathbf{u}}_i; \hat{\mathbf{u}}'_i, \Sigma'_i) \ell\left(\mathbf{y}_i, \tilde{\mathbf{u}}_i; \boldsymbol{\theta}\right) d\tilde{\mathbf{u}}_i$, where Φ is the multivariate normal probability density function with mean

$$\hat{\mathbf{u}}'_i = \left(\sum_j Z_{ij}'^T R'^{-1} Z'_{ij} + V'^{-1} \right)^{-1} \left(\sum_j Z_{ij}'^T R'^{-1} (\mathbf{y}_{ij} - \mathbf{b}') + V'^{-1} \mathbf{m}' \right), \quad (10)$$

and covariance $\Sigma'_i = \left(\sum_j Z_{ij}'^T R'^{-1} Z'_{ij} + V'^{-1} \right)^{-1}$. Note that for individuals with a single visit, $\sum_j Z_{ij}'^T R'^{-1} Z'_{ij}$ is a singular matrix. Considering $\mathbf{u}_i$ as parameters, as in Jedynak et al. (2012) or Bilgel et al. (2015c), is equivalent to assuming that V has infinitely large diagonal elements (i.e., an uninformative uniform prior) such that its inverse disappears. Therefore, it is not possible to compute $\hat{\mathbf{u}}'_i$ for individuals having only one visit using this approach. On the other hand, incorporation of a bivariate normal prior on the subject-specific variables $\mathbf{u}_i$ allows Eq. 10 to be computed for individuals with a single visit.

Evaluation of the E-step integral involves second moments of a Gaussian random variable. We ignore the terms that do not depend on $\boldsymbol{\theta}$ as they will not

be relevant in the maximization step and obtain:

$$\begin{aligned} Q\left(\boldsymbol{\theta},\boldsymbol{\theta}'\right) &= -\frac{1}{2}\sum_{i,j}\log|R| \\ &\quad -\frac{1}{2}\sum_{i,j}(\mathbf{y}_{ij}-Z_{ij}\hat{\mathbf{u}}_i'-\mathbf{b})^T R^{-1}(\mathbf{y}_{ij}-Z_{ij}\hat{\mathbf{u}}_i'-\mathbf{b}) \\ &\quad -\frac{1}{2}\sum_{i,j}\mathrm{Tr}\left(Z_{ij}^T R^{-1} Z_{ij}\Sigma_i'\right)-\frac{1}{2}\sum_i \log|V| \\ &\quad -\frac{1}{2}\sum_i(\hat{\mathbf{u}}_i'-\mathbf{m})^T V^{-1}(\hat{\mathbf{u}}_i'-\mathbf{m})-\frac{1}{2}\sum_i \mathrm{Tr}\left(V^{-1}\Sigma_i'\right). \quad (11) \end{aligned}$$

2.3. M-step

Here, we provide the update equations for the EM algorithm obtained by maximizing $Q(\boldsymbol{\theta},\boldsymbol{\theta}')$ with respect to each parameter, and provide derivations for these update equations in the Appendix. The update equations depend on previous parameter estimates $\boldsymbol{\theta}' = \{\mathbf{m}', \boldsymbol{\nu}', \mathbf{a}', \mathbf{b}', \boldsymbol{\lambda}', \boldsymbol{\rho}'\}$ as well as the progression score estimates $s_{ij}' = \mathbf{q}_{ij}^T\hat{\mathbf{u}}_i'$, where $\hat{\mathbf{u}}_i'$ is as given in Eq. 10:

$$\mathbf{a} = \frac{\left(\sum_i v_i\right)\left(\sum_{i,j}\mathbf{y}_{ij}s_{ij}'\right)-\left(\sum_{i,j}\mathbf{y}_{ij}\right)\left(\sum_{i,j}s_{ij}'\right)}{\left(\sum_i v_i\right)\left(\sum_{i,j}\mathbf{q}_{ij}^T\Sigma_i'\mathbf{q}_{ij}+s_{ij}'^2\right)-\left(\sum_{i,j}s_{ij}'\right)^2}, \quad (12)$$

$$\mathbf{b} = \frac{\left(\sum_{i,j}\mathbf{y}_{ij}\right)\left(\sum_{i,j}\mathbf{q}_{ij}^T\Sigma_i'\mathbf{q}_{ij}+s_{ij}'^2\right)-\left(\sum_{i,j}\mathbf{y}_{ij}s_{ij}'\right)\left(\sum_{i,j}s_{ij}'\right)}{\left(\sum_i v_i\right)\left(\sum_{i,j}\mathbf{q}_{ij}^T\Sigma_i'\mathbf{q}_{ij}+s_{ij}'^2\right)-\left(\sum_{i,j}s_{ij}'\right)^2}, \quad (13)$$

$$\mathbf{m} = \frac{1}{n}\sum_i \hat{\mathbf{u}}_i', \quad (14)$$

$$\boldsymbol{\nu} = \arg\max_{\boldsymbol{\nu}} Q(\boldsymbol{\theta},\boldsymbol{\theta}'), \quad (15)$$

$$\boldsymbol{\lambda},\,\boldsymbol{\rho} = \arg\max_{\boldsymbol{\lambda},\boldsymbol{\rho}} Q(\boldsymbol{\theta},\boldsymbol{\theta}'). \quad (16)$$

Note that if C is fixed to be the identity matrix, a closed form solution for $\boldsymbol{\lambda}$ exists, as given in Equation A.15. Once the optimal parameters are found, the subject-specific variables are predicted using Eq. (10).

2.4. Parameter standardization

As described in Proposition A.2, there are certain reparameterizations that yield identical models. For example, one can multiply the trajectory slope

parameters by 2 and divide all progression scores by 2 (which is achieved by dividing all α and β values by 2) without altering the model. This is the scaling degree of freedom. There is also a translation degree of freedom. We account for these degrees of freedom and anchor the model by calibrating the progression score scale. We calibrate such that baseline PS has a mean of 0 and a standard deviation of 1. This involves replacing the model parameters $\{\mathbf{m}, V(\boldsymbol{\nu}), \mathbf{a}, \mathbf{b}, R(\boldsymbol{\lambda}, \boldsymbol{\rho})\}$ with $\left\{ w\mathbf{m} + \begin{bmatrix} 0 \\ z \end{bmatrix}, w^2 V(\boldsymbol{\nu}), \frac{1}{w}\mathbf{a}, \mathbf{b} - \frac{z}{w}\mathbf{a}, R(\boldsymbol{\lambda}, \boldsymbol{\rho}) \right\}$, where $-\frac{z}{w} = \frac{1}{n}\sum_i s_{i1}$ is the mean progression score at baseline, and $\frac{1}{w} = \sqrt{\frac{1}{n}\sum_i \left(s_{i1} - \frac{1}{n}\sum_i s_{i1}\right)^2}$ is the standard deviation of baseline progression scores. We standardize the subject-specific estimates α_i, β_i, and s_{ij} accordingly: $\alpha_i^* = w\alpha_i, \beta_i^* = w\beta_i + z$, and $s_{ij}^* = ws_{ij} + z$. This reparametrization yields a PS scale where 0 corresponds to the sample average at baseline and the variance of PS at baseline is 1.

2.5. Implementation details

We first fit the model assuming that $C = I_{K\times K}$ and Λ is a diagonal matrix with positive elements $\boldsymbol{\lambda}$. We denote the estimated Λ in this model as $\hat{\Lambda}$. The estimates obtained from this model for the parameters $\mathbf{a}, \mathbf{b}, \mathbf{m}, V$ are used as initializations in the model where $C = C(\rho)$. In this model where correlations are taken into account, we assume that $\Lambda = \lambda\hat{\Lambda}$, where λ is an unknown parameter to be estimated. The spatial correlation function among those presented in Table 1 that results in the highest log-likelihood value is chosen for the final model.

The EM algorithm is implemented in MATLAB 8.1 and Statistics Toolbox 8.2 (The MathWorks Inc., Natick, MA). Our code is freely available online.[1]

2.6. Confidence intervals

We use bootstrapping via Monte Carlo resampling to estimate confidence intervals for each model parameter. We sample with replacement from the

[1] `https://www.iacl.ece.jhu.edu/Resources`

original collection of subjects to generate a new dataset containing an equal number of subjects and fit the model on this generated sample. This sampling and fitting procedure is repeated to generate bootstrap estimates. We then compute 95% confidence intervals for each parameter across the bootstrap estimates. In the bootstrap experiments, we fix the value of ρ at its estimate on the whole sample to enable faster computation.

2.7. Comparison to linear mixed effects model

We compared our model to a linear mixed effects (LME) model that included random intercepts and slopes at each voxel. The LME model for subject i, visit j and voxel k is given by

$$y_{ijk} = (\eta_k + \eta_{ik})t_{ij} + (\gamma_k + \gamma_{ik}) + \varepsilon_{ijk}, \tag{17}$$

where $\begin{pmatrix}\eta_k \\ \gamma_k\end{pmatrix}$ are the fixed effects, $\begin{pmatrix}\eta_{ik} \\ \gamma_{ik}\end{pmatrix} \sim \mathcal{N}(0, \Xi_k)$ are the random effects and $\varepsilon_{ijk} \sim \mathcal{N}(0, \sigma_k^2)$ is the observation noise. We used the LME implementation in MATLAB Statistics Toolbox 8.2.

2.8. Simulated data set

We simulated visits such that the sample was similar to our PET data in terms of number of visits per subjects and age range. We generated a data set with 100 individuals, each with up to 7 visits with $5 \times 5 \times 5$ images with 4 mm isotropic voxels. We fixed the ground truth values of the model parameters $\boldsymbol{\theta} = \{\mathbf{m}, \boldsymbol{\nu}, \mathbf{a}, \mathbf{b}, \boldsymbol{\lambda}, \rho\}$ at values close to those we observed in exploratory models fitted to DVR data. We generated $\mathbf{u}_i$ from a bivariate normal distribution with mean $\mathbf{m}$ and variance $V(\boldsymbol{\nu})$. The progression score for each visit was then computed as $s_{ij} = \mathbf{q}_{ij}^T\mathbf{u}$, and observations were generated using the PS model. We performed 1000 bootstrap iterations to obtain 95% confidence intervals for each parameter and subject-specific variable. We computed the cosine similarity for each variable for each bootstrap experiment using

$$\frac{\boldsymbol{\phi}^T\widehat{\boldsymbol{\phi}}}{\|\boldsymbol{\phi}\|_2\|\widehat{\boldsymbol{\phi}}\|_2}, \tag{18}$$

where $\widehat{\boldsymbol{\phi}}$ is the estimate of the variable of interest ($\mathbf{a}, \mathbf{b}, \boldsymbol{\alpha}, \boldsymbol{\beta}$, or $\mathbf{s}$) and $\boldsymbol{\phi}$ is the corresponding ground truth value. A value of 1 indicates a perfect estimate.

2.9. Amyloid imaging data set

We used longitudinal positron emission tomography (PET) data for participants from the Baltimore Longitudinal Study of Aging (Shock et al., 1984) neuroimaging substudy (Resnick et al., 2000). Participant demographics are presented in Table 2. Starting in 2005, amyloid-beta (Aβ) PET scans were acquired on a GE Advance scanner over 70 minutes following an intravenous bolus injection of Pittsburgh compound B (PiB), yielding dynamic PET scans with 33 time frames. PET images were reconstructed using filtered backprojection with a ramp filter, yielding a spatial resolution of approximately 4.5 mm full width at half max at the center of the field of view (image matrix = 128×128, 35 slices, pixel size = 2×2 mm, slice thickness = 4.25 mm).

The frames of each dynamic PiB-PET scan were aligned to the average of the first two minutes to remove motion (Jenkinson et al., 2002). For registration purposes, we obtained static images by averaging the first 20 minutes of the dynamic PiB-PET scan. Follow-up PiB-PET scans were rigidly registered onto the baseline PiB-PET within each participant using the 20-minute average images. Baseline magnetic resonance images (MRIs) were rigidly registered onto their corresponding 20-minute PiB-PET average, and their FreeSurfer segmentations (Dale et al., 1999; Desikan et al., 2006) were transformed accordingly. Distribution volume ratio (DVR) images were calculated in the native space of each PiB-PET image using the simplified reference tissue model with the cerebellar gray matter as reference tissue (Zhou et al., 2003). The MRIs coregistered with the PET were deformably registered (Avants et al., 2008) onto a study-specific template (Avants et al., 2010; Bilgel et al., 2015b) and transformed to 4 mm isotropic MNI space using a pre-calculated affine transformation. The resulting mappings were applied to the DVR images that have been registered to baseline to bring them into the MNI space. We used all voxels within the brain mask in the MNI space to fit the PS model.

Table 2: Participant demographics. MCI = mild cognitive impairment, SD = standard deviation.

Characteristic	$N = 104$
Baseline age in years, mean (SD)	77.0 (7.9)
Range	55.7–93.4
Female, n (%)	48 (46%)
PiB-PET scans, n	300
PiB-PET per subject, n	2.9 (1.8)
Range	1–7
Years between first and last scan, mean (SD)	3.3 (2.9)
Range	0.0–9.0
Mean cortical DVR at last visit (SD)	1.12 (0.19)
PiB+ at last visit, n (%)	42 (40%)
MCI diagnosis at last visit, n (%)	3 (3%)
Dementia diagnosis at last visit, n (%)	4 (4%)

Mean cortical DVR is a widely used measure obtained from PiB-PET images for quantifying the level of brain amyloid. We computed mean cortical DVR for each PiB-PET image by averaging the voxelwise DVR values across cingulate, frontal, parietal, lateral temporal, and lateral occipital cortices, excluding the sensorimotor strip. We used a mean cortical DVR threshold of 1.06, which was computed from a two-class Gaussian mixture model fitted on baseline measures, to separate individuals into PiB- and PiB+ groups, as described in Bilgel et al. (2015a).

2.10. Hypothesis testing

The confidence intervals obtained via bootstrapping allow for hypothesis testing. For the amyloid images, we focus on studying the precuneus, since previous cross-sectional amyloid imaging studies have suggested that precuneus has the highest deposition levels (Mintun et al., 2006) and provided preliminary

evidence that it may be the most rapid accumulator (Rodrigue et al., 2012) among cortical regions. Our specific hypotheses are as follows:

1. The precuneus has the highest amyloid load along stages of amyloid accumulation.
2. The precuneus accumulates amyloid faster than other cortical regions.

We refer to the progression scores calculated using the DVR images as Aβ-PS. To test the first hypothesis, we compare the amyloid levels in the precuneus with other regions at various Aβ-PS values spanning the range observed in our data set. Our purpose in performing these comparisons is to shed light onto the temporal ordering of changes in different cortical regions. For each bootstrap experiment, we average the predicted DVR levels $\hat{y}_k(s) = a_k s + b_k$ within each ROI to obtain

$$\hat{y}_r(s) = \frac{1}{|\text{ROI}_r|} \sum_{k \in \text{ROI}_r} \hat{y}_k(s), \tag{19}$$

where k is the voxel index and r is the ROI index. We then compute the following statistic for each bootstrap:

$$T = \hat{y}_{\text{precuneus}}(s) - \max_{r \neq \text{precuneus}} \hat{y}_r(s). \tag{20}$$

We reject the null hypothesis that the amyloid load in the precuneus at Aβ-PS $= s$ is not different from other regions at significance level γ if the two-sided $100(1-\gamma)$ confidence interval of the test statistic T, computed using the bootstrap estimates of T, does not contain 0. The smallest value of γ such that the two-sided $100(1-\gamma)$ confidence interval of the test statistic T contains 0 is the p-value of the test.

To test the second hypothesis, we pursue a similar approach using the trajectory slope parameter $\mathbf{a}$. For each bootstrap experiment, we obtain the average a_k per cortical ROI to obtain a_r, where r the ROI index. We then compute the following statistic for each bootstrap:

$$T = a_{\text{precuneus}} - \max_{r \neq \text{precuneus}} a_r. \tag{21}$$

The condition for the rejection of the null hypothesis that the rate of amyloid accumulation in the precuneus is not different from that of other regions is as described previously.

3. Results

3.1. Simulation

The Akaike information criterion (AIC) was 39.0×10^3 for the LME model, 32.0×10^3 for the PS model where $C = I_{K\times K}$, and -1.6×10^3 for the PS model where $C = C(\rho)$, indicating that the PS model where spatial correlations are modeled fits the data the best. We computed the percentage of variables that are "correct" by counting the variables whose ground truth values fell within the 95% confidence interval computed via bootstrapping. For example, for $\mathbf{a}$ and $\mathbf{b}$, a value of 10% indicates that 10% of the biomarkers had a correct estimate for a_k and b_k. For $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$, 10% indicates that 10% of the individuals had a correct estimate for α_i and β_i. For $\mathbf{s}$, 10% indicates that 10% of the visits had a correct estimate for s_{ij}. Simulation results are presented in Table 3. When there are correlations in the data, not modeling them yields inaccurate estimates for the trajectory slope parameter $\mathbf{a}$, the subject-specific variables α_i, β_i, and the progression scores s_{ij}. Using the correlation model improves these estimates significantly.

3.2. Amyloid images

In our preliminary analysis of the noise spatial correlation structure using the semivariogram (Cressie and Hawkins, 1980; Bilgel et al., 2015c), we observed that the rational quadratic had the best fit to the empirical semivariogram (Inline Supplementary Fig. 1). Rational quadratic also yielded the fit with largest log-likelihood, and thus was selected as the correlation function for the model. AIC was -1.2×10^7 for the LME model, -1.1×10^7 for the model where $C = I_{K\times K}$, and -2.0×10^7 for the model where $C = C(\rho)$, indicating that the PS model where spatial correlations are modeled fits the data the best. The

Table 3: Simulation results. Mean cosine similarity values across 1000 bootstrap experiments (with standard deviation) and percentage of variable elements that are correct based on 95% confidence intervals are presented.

	$C = I_{K\times K}$		$C = C(\rho)$	
Variable	Cosine similarity	% correct	Cosine similarity	% correct
$\mathbf{a}$	0.9514 ± 0.0154	24	0.9821 ± 0.0073	98
$\mathbf{b}$	0.9998 ± 0.0001	99	0.9998 ± 0.0001	98
$\boldsymbol{\alpha}$	0.1587 ± 0.0812	18	0.7922 ± 0.0332	40
$\boldsymbol{\beta}$	0.1250 ± 0.0794	19	0.7835 ± 0.0339	38
$\mathbf{s}$	0.8017 ± 0.0443	23	0.9881 ± 0.0053	85

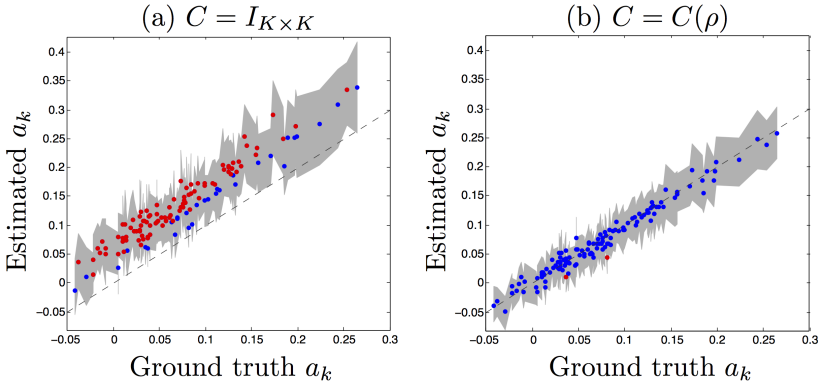

Figure 3: Estimated trajectory slope parameters a_k vs ground truth. Dashed line indicates $x = y$. Gray band corresponds to the 95% confidence intervals obtained from bootstrapping. Estimates are shown in blue if their 95% confidence interval intersects the $x = y$ line, and in red otherwise. Results from the model where (a) $C = I_{K\times K}$ and (b) $C = C(\rho)$.

spatial correlation parameter ρ for the rational quadratic correlation model was estimated to be 4.5 mm, and λ was estimated to be 0.929. A comparison of the PS values obtained from the different correlation models are presented in the Supplementary Material (Inline Supplementary Fig. 2). Model fitting using 104 subjects with a total of 300 visits each having about $30,000$ brain voxels took approximately 5 seconds on an Intel Xeon 8-core 3.1 GHz machine with 128 GB RAM for the case where voxels are assumed to be independent (i.e., $C = I_{K\times K}$).

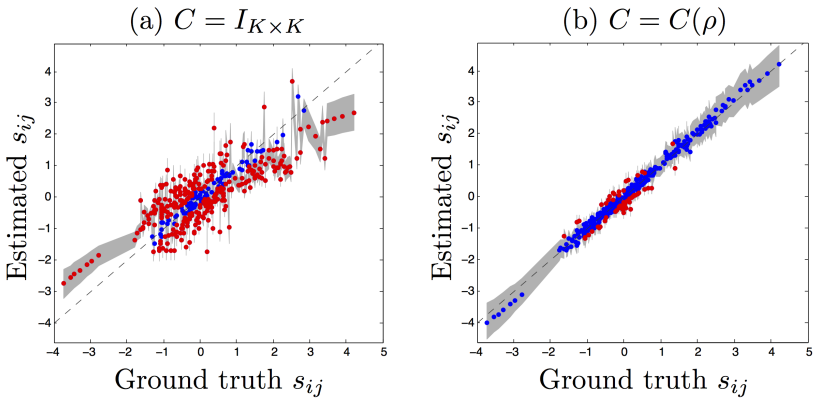

Figure 4: Predicted progression scores s_{ij} vs ground truth. Dashed line indicates $x = y$. Gray band corresponds to the 95% confidence intervals obtained from bootstrapping. Estimates are shown in blue if their 95% confidence interval intersects the $x = y$ line, and in red otherwise. Results from the model where (a) $C = I_{K\times K}$ and (b) $C = C(\rho)$.

The following model fitting with $C = C(\rho)$ took approximately 30 minutes per EM iteration, with convergence achieved in 2 to 6 iterations depending on the correlation structure.

3.2.1. Comparison of Aβ-PS to mean cortical DVR

We compared the subject-specific variables obtained from model fitting to empirical values obtained from mean cortical DVR, which is the widely used measure for quantifying levels of brain amyloid (Fig. 5). For each individual with at least two visits, we fit a line to the mean cortical DVR data to estimate the slope of amyloid accumulation as well as the intercept. The subject-specific variable α_i, which represents the rate of amyloid progression, explained 62% of the variability in the empirical slope of amyloid accumulation computed from longitudinal mean cortical DVR. We observed a much higher correlation between mean cortical DVR and Aβ-PS ($R^2 = 0.95$ at baseline and 0.96 at last visit).

When plotted against age, mean cortical DVR and Aβ-PS revealed similar patterns (Fig. 6). The progression as measured using voxelwise DVR data is not linearly associated with age, and the Aβ-PS was able to capture this.

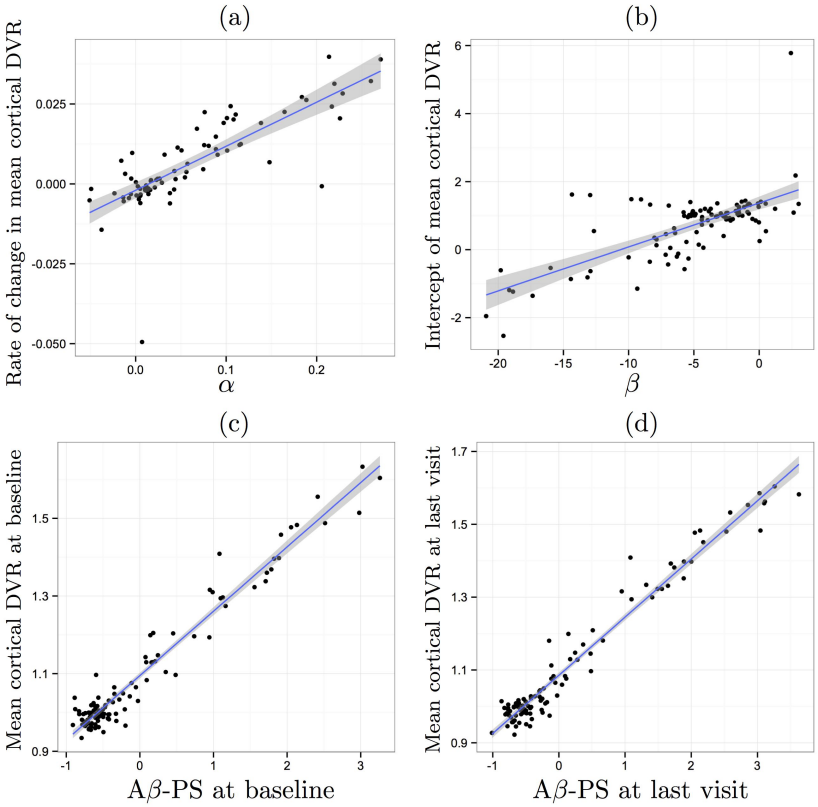

Figure 5: Correlation of estimated subject-specific variables with mean cortical DVR measures. The line of best fit is shown in blue, and its 95% confidence band in gray. (a) Rate of annual change in mean cortical DVR vs. α, the predicted rate of change in amyloid progression score ($R^2 = 0.62$). (b) Intercept of mean cortical DVR vs. β, the progression score intercept ($R^2 = 0.48$). (c) Mean cortical DVR vs. Aβ-PS at baseline ($R^2 = 0.95$). (d) Mean cortical DVR vs. Aβ-PS at last visit ($R^2 = 0.96$).

3.2.2. Amyloid trajectories

The trajectory slope parameters a_k obtained from the PS model (Fig. 7) revealed symmetric patterns across the cerebral hemispheres. The precuneus and frontal lobe showed the greatest increases in DVR with Aβ-PS, smaller increases in lateral temporal and temporoparietal regions, and minimal increases in the occipital lobe and the sensorimotor strip. Voxelwise trajectories are further illustrated in Fig. 8, which shows the predicted DVR values on the cortical

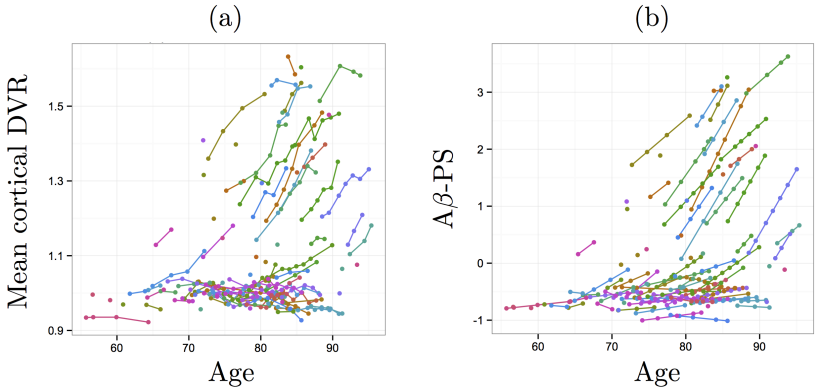

Figure 6: (a) Mean cortical DVR and (b) Aβ-PS plotted against age. Longitudinal data points are connected by lines within each subject. Different colors indicate different subjects.

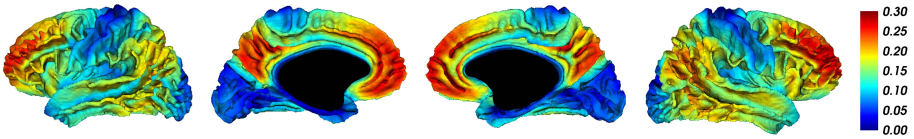

Figure 7: Slope parameters a_k obtained from voxelwise PS model projected onto the cortical surface. For each unit increase in Aβ-PS, the DVR value at voxel k increases by a_k.

surface at three Aβ-PS levels.

In order to better investigate regional trends, we averaged the PS model results within each cortical ROI and plotted these ROI averages as a function of Aβ-PS (Fig. 9). We used bootstrap results to compute 95% confidence bands for these ROI averages. Based on the 95% confidence band of the precuneus illustrated in Fig. 9, precuneus appears to be the earliest accumulator and has the highest amyloid levels through late stages of amyloid accumulation (Aβ-PS$\geq$ 2). We present confidence bands for other cortical regions in Fig. 10.

Based on our hypothesis testing procedure, we found that precuneus had the highest amyloid levels at Aβ-PS values of -0.5, 0, 1, 2, and 3 (all $p < 0.01$). On the other hand, while the estimated rate of amyloid accumulation was highest in the precuneus, this was not statistically significant ($p = 0.33$). A comparison of the levels of amyloid at Aβ-PS= 0 and rates of accumulation across ROIs is

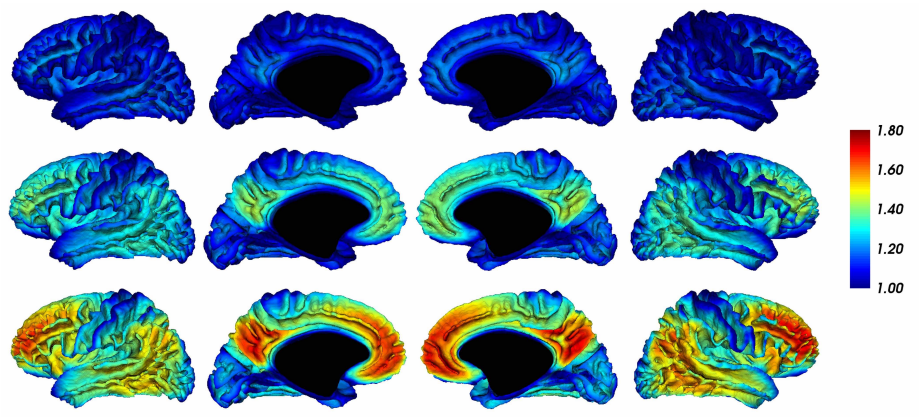

Figure 8: Predicted DVR levels at Aβ-PS $= -0.6$ (*top*), 0.4 (*middle row*), and 1.5 (*bottom*).

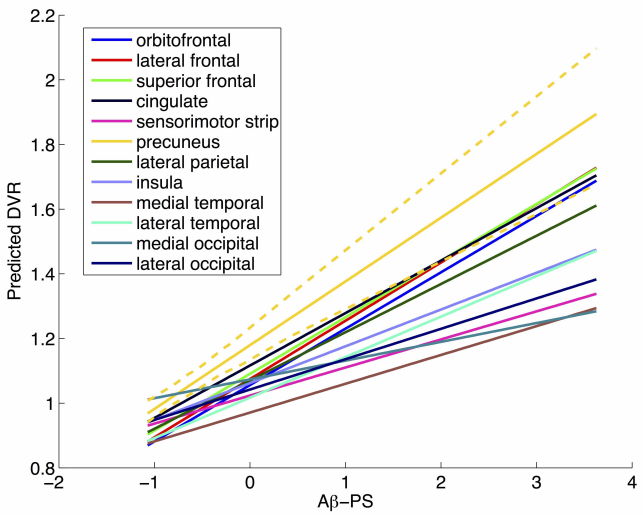

Figure 9: Regional trajectories as function of Aβ-PS. The PS model was used to make voxelwise predictions at a range of Aβ-PS values, and these predictions were averaged within each ROI to obtain regional trajectories. The dashed lines indicate the 95% confidence band obtained using bootstrap results for the precuneus.

presented in Fig. 11.

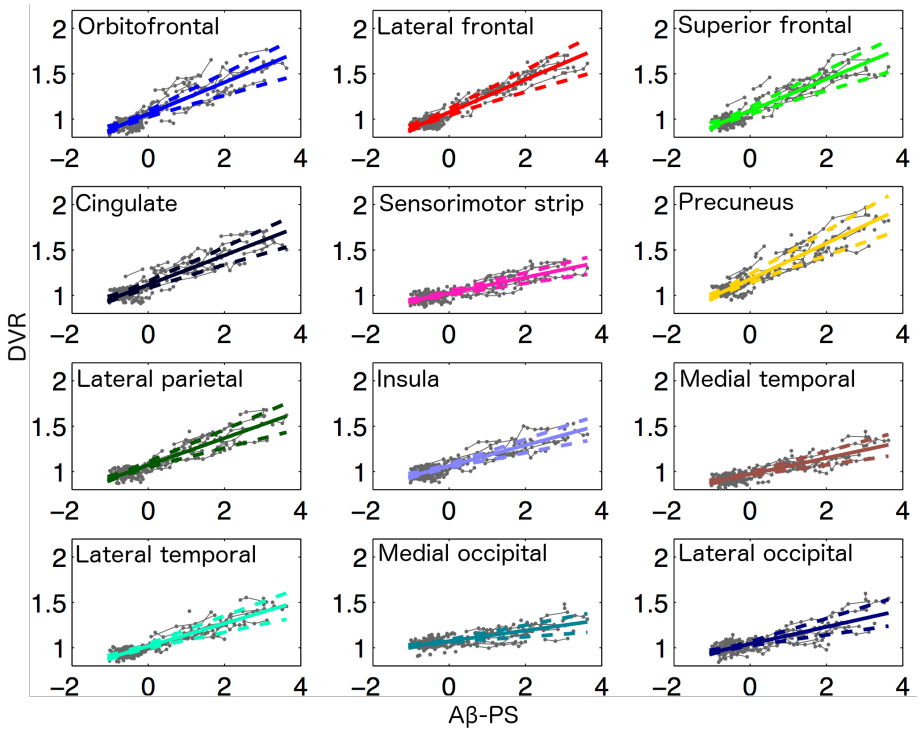

Figure 10: Regional trajectories as function of Aβ-PS. The PS model was used to make voxelwise predictions at a range of Aβ-PS values, and these predictions were averaged within each ROI to obtain regional trajectories. The dashed lines indicate the 95% confidence bands for the cortical regions. Estimated trajectories with their 95% confidence bands are superimposed on observed longitudinal data (in gray).

3.2.3. Comparison to LME results

The trajectories we obtained using the PS model were consistent with the results of the LME model (Inline Supplementary Figs. 3, 4). However, the fixed effects obtained from the LME model did not describe the observed data as well as the trajectory slopes obtained from the PS model (Inline Supplementary Fig. 5). Using the LME model estimates and a bootstrapping procedure similar to the one we applied for the PS model, we found that the precuneus had the highest regional amyloid levels compared to other cortical regions at ages 70, 80, and 90 (all $p < 0.01$), but not at age 65 ($p = 0.12$). Similar to our findings based on the PS model, the rate of amyloid accumulation was highest in the precuneus,

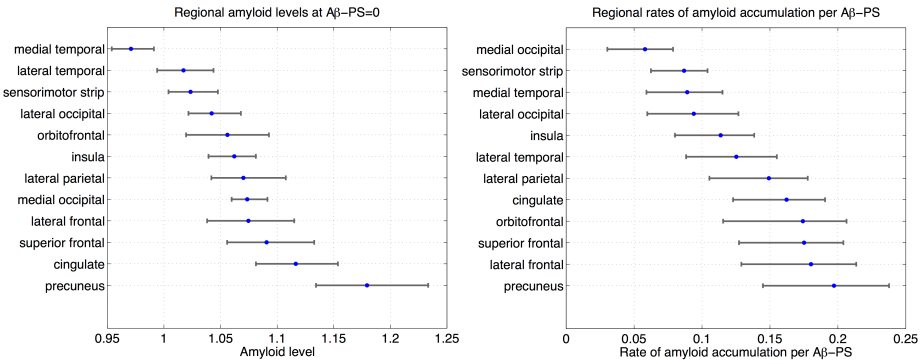

Figure 11: Comparison of levels of amyloid at Aβ-PS$=0$ and rates of amyloid accumulation across cortical regions. The intercept parameter **b** obtained from the PS model was averaged within each ROI to obtain the regional amyloid levels at Aβ-PS$=0$, and the trajectory slope parameter **a** was averaged within each ROI to obtain regional rates.

but this was not statistically significant ($p = 0.75$). A comparison of the levels of amyloid at the mean baseline age of the sample and rates of accumulation across ROIs is presented in Inline Supplementary Fig. 6.

4. Discussion

We presented a statistical model for estimating longitudinal trajectories of voxelwise neuroimaging data. Our model is based on the concept that age is not a good metric for disease progression since each individual has his or her own onset and rate of progression of disease. We accounted for these inter-individual differences to temporally align individuals based on their collection of voxelwise measurements. As a result of this temporal alignment, we obtained a metric of disease or underlying biological process stage as reflected in the neuroimaging data, and we call this metric "progression score". By analyzing voxelwise neuroimaging data as a function of progression score rather than time, we constructed trajectories that better represent changes occurring with disease stage.

Simulation results showed that the progression score model fits the data better than the linear mixed effects model, as evidenced by AIC. Furthermore,

the results showed that modeling spatial correlations is important for extracting accurate summary scores from data with underlying correlations. In the absence of correlation modeling when the underlying data are correlated, estimates of the trajectory slopes $\mathbf{a}$ as well as subject-specific variables α, β, and s were adversely affected. The difference in the subject-specific variables was mainly due to inaccurate estimates of the prior covariance V. Note that

$$\text{Cov}(\mathbf{y}_{ij}) = \mathbf{a}\mathbf{q}_{ij}^T V \mathbf{q}_{ij}\mathbf{a}^T + R, \tag{22}$$

which explains why both $\mathbf{a}$ and V are affected when correlations are not modeled through the covariance matrix R. This highlights the importance of modeling spatial correlations in data where such effects are evident, especially in order to estimate correct trajectory slopes and individualized progression scores. To further understand this phenomenon, we conducted simulations (data not presented) where we duplicated biomarkers with lower signal-to-noise ratios and repeated model fitting including these duplicate biomarkers. We observed that the parameter estimates as well as the PS values were biased if we assumed independence across biomarkers. This might be because the model interprets each of the duplicated biomarkers as a separate piece of evidence towards the computation of individual progression scores. However, when we included a proper noise correlation model, we were able to recover our original results on the collection of unique biomarkers.

The proposed EM framework simplifies the estimation of spatial correlation parameters, and accurately estimates the trajectory parameters $\mathbf{a}$ and $\mathbf{b}$ as well as predicting the progression scores s. The performance of the fitting procedure in predicting α and β is worse than for s. This may be due to the modeling of α and β as random variables that do not directly contribute to the observations $\mathbf{y}$. We observed a similar pattern in the amyloid data: the agreement between estimated Aβ-PS and mean cortical DVR was greater than the agreement between estimated α and the rate of change of mean cortical DVR per individual. This suggests that the progression scores are more reliable than the subject-specific variables. Aβ-PS was highly correlated with mean cortical DVR, a widely used

measure for quantifying PiB-PET scans and assessing longitudinal change. This indicates that Aβ-PS is a meaningful score extracted from voxelwise imaging data. Unlike mean cortical DVR, there are no *a priori* assumptions regarding which regions or voxels should be included in the computation of Aβ-PS. This property of PS allows for the discovery of new patterns in longitudinal data.

Trajectories obtained from the PS model captured a wider dynamic range of DVR values and had a better fit to the observed data compared to the trajectories obtained using the LME model fixed effects. This difference between the PS and LME models underscores the fact that age is not an appropriate metric for staging amyloid accumulation. By aligning individuals based on their amyloid scans, the PS model enables a better temporal metric of amyloid staging.

Aβ-PS allows for the exploration of longitudinal voxelwise trajectories within a hypothesis testing framework due to its underlying statistical model. Our results suggest that the precuneus exhibits the earliest cortical amyloid changes, but that its rate of amyloid accumulation does not differ significantly from other cortical regions. Based on a qualitative evaluation of our estimated trajectories, we found that amyloid accumulation in the precuneus is followed by cingulate and frontal cortices, then by lateral parietal cortex, followed by insula and lateral temporal cortex. We observed minimal amyloid accumulation in visual cortex, hippocampal formation and the sensorimotor strip, which agrees with previous reports that these regions accumulate amyloid in later stages of AD (Braak and Braak, 1991). Previous reports have highlighted precuneus as an early amyloid accumulator (Mintun et al., 2006; Rodrigue et al., 2012) as well as frontal, cingulate, and parietal regions (Jack et al., 2008). Contrary to our finding highlighting precuneus as the earliest cortical accumulator, a study of cognitively normal adults found that the right medial frontal cortex accumulates amyloid the earliest, closely followed by bilateral precuneus based on cross-sectional voxelwise analyses (Villeneuve et al., 2015). We presented regional trajectories averaged bilaterally in this work; however, hemisphere-specific trajectories were consistent with our bilaterally-averaged trajectories and we did not observe that medial frontal cortex precedes precuneus. Further studies with larger samples will be

instrumental in elucidating the regional progression of amyloid accumulation.

One of the strengths of our model for the progression scores s is that there are no *a priori* assumptions on the global form of $s(t)$, which can be thought of as an underlying function that describes the evolution of progression score over time. Instead, the model makes linear local approximations to $s(t)$ per individual. The linear form we use makes the EM approach analytically tractable and enables the discovery of the global form of $s(t)$. On the other hand, the linear relationship we assume between s and t is also a limitation since it may be an oversimplification over long follow-up durations. In order to capture dynamics over longer periods more accurately, it is necessary to investigate the relationship between the progression scores s and time t, and select an appropriate function to link these variables.

Another limitation of our model is the assumption of linear biomarker trajectories. This assumption is not reflective of the fact that voxelwise DVR has a theoretical minimum of 1 (while in practice, there are DVR values lower than 1 due to noise). Furthermore, several studies of longitudinal amyloid deposition have found evidence for a ceiling effect in amyloid deposition in late AD, suggesting a sigmoid trajectory for amyloid levels (Villemagne et al., 2013; Villain et al., 2012). Our use of a linear rather than sigmoid trajectory inevitably results in inaccuracies in PS calculation for individuals whose voxelwise data lie along the plateaus of the sigmoid in reality. The linear trajectory assumption may also prevent us from characterizing subtler differences in the temporal ordering of the onset of amyloid accumulation across different regions. However, linear trajectories, unlike sigmoid trajectories, yield closed-form update equations for trajectory parameters, greatly facilitating the model fitting procedure.

The spatial covariance functions we used in our analyses are covariance functions of strictly stationary isotropic processes. The underlying spatial noise process is affected by the PET scanner, image reconstruction algorithm, radiotracer delivery and binding in different brain regions, kinetic parameter estimation algorithm, and registration methods used to bring all scans in alignment. To accurately model these influences on noise, complex noise spatial

covariance models are needed. In this work, we made simplifying assumptions on the noise properties to facilitate the study of longitudinal trajectories of amyloid images. As a result of these simplifications, our model does not accurately capture noise properties; however, by eliminating the inaccurate assumption of independence across voxels, our approach improves the accuracy of the estimated trajectories and progression scores. Refined spatial noise models incorporating non-stationarity and anisotropy may further improve the estimation accuracy.

Our model can be applied to studying higher resolution images. To reduce computational memory burden, it may be necessary to impose sparsity on the spatial correlation matrix, which can be done by imposing a correlation value of 0 at large distances. The sparsity property of the correlation matrix can be taken advantage of to yield Cholesky decompositions using less memory and time.

In conclusion, the progression score model allows for extracting summary scores from longitudinal data for each scan. In this work, we have extended the progression score model proposed by Jedynak et al. (2012) to voxelwise imaging data by accounting for spatial correlations and enabling efficient handling of the large number of voxels through the EM framework. The incorporation of a prior on subject-specific variables allowed for the inclusion of individuals with a single visit in the model. Our method can be extended to the analysis of other types of imaging data, and in cases where a summary score such as mean cortical DVR is not available, the progression score estimates can be highly informative for progression staging.

5. Acknowledgment

We thank Kalyani Kansal for her thorough review of the mathematical derivations and assistance with code testing. This research was supported in part by the Intramural Research Program of the National Institutes of Health and by the Michael J. Fox Foundation for Parkinson's Research, MJFF Research Grant ID: 9310.

Appendix A.

Proposition A.1.

$$f(\tilde{\mathbf{u}}_i \mid \mathbf{y}_i; \boldsymbol{\theta}') \quad \propto \quad \Phi(\tilde{\mathbf{u}}_i; \hat{\mathbf{u}}_i', \Sigma_i'), \tag{A.1}$$

where $\Phi(\,\cdot\,; \boldsymbol{\mu}, \Sigma)$ *denotes the probability density function of a multivariate normal with mean* $\boldsymbol{\mu}$ *and covariance matrix* Σ*,*

$$\hat{\mathbf{u}}_i' = \left(\sum_j Z_{ij}'^T R'^{-1} Z_{ij}' + V'^{-1} \right)^{-1} \left(\sum_j Z_{ij}'^T R'^{-1} (\mathbf{y}_{ij} - \mathbf{b}') + V'^{-1} \mathbf{m}' \right), \tag{A.2}$$

and $\Sigma_i' = \left(\sum_j Z_{ij}'^T R'^{-1} Z_{ij}' + V'^{-1} \right)^{-1}$ *is a covariance matrix.*

Proof. Independence across visits allows us to write

$$\begin{aligned} f(\mathbf{y}_i \mid \tilde{\mathbf{u}}_i; \boldsymbol{\theta}') &= \prod_j f(\mathbf{y}_{ij} \mid \tilde{\mathbf{u}}_i; \boldsymbol{\theta}') \\ &= \prod_j \Phi(\mathbf{y}_{ij}; Z_{ij}' \tilde{\mathbf{u}}_i + \mathbf{b}', R') \\ &\propto \exp\left\{ -\frac{1}{2} \left[\tilde{\mathbf{u}}_i^T \left(\sum_j Z_{ij}'^T R'^{-1} Z_{ij}' \right) \tilde{\mathbf{u}}_i \right. \right. \\ &\qquad \left. \left. -2 \left(\sum_j (\mathbf{y}_{ij} - \mathbf{b}')^T R'^{-1} Z_{ij}' \right) \tilde{\mathbf{u}}_i \right] \right\}, \end{aligned} \tag{A.3}$$

where we have ignored the terms that do not depend on $\tilde{\mathbf{u}}_i$.

By Bayes' rule,

$$f(\tilde{\mathbf{u}}_i \mid \mathbf{y}_i; \boldsymbol{\theta}') \quad \propto \quad f(\mathbf{y}_i \mid \tilde{\mathbf{u}}_i; \boldsymbol{\theta}') f(\tilde{\mathbf{u}}_i; \boldsymbol{\theta}') \tag{A.4}$$

$$\propto \quad \exp\left\{ -\frac{1}{2} \left[\tilde{\mathbf{u}}_i^T \left(\sum_j Z_{ij}'^T R'^{-1} Z_{ij}' + V'^{-1} \right) \tilde{\mathbf{u}}_i \right. \right.$$

$$\left. \left. -2 \left(\sum_j (\mathbf{y}_{ij} - \mathbf{b}')^T R'^{-1} Z_{ij}' + \mathbf{m}'^T V'^{-1} \right) \tilde{\mathbf{u}}_i \right] \right\} \tag{A.5}$$

$$\propto \quad \Phi(\tilde{\mathbf{u}}_i; \hat{\mathbf{u}}_i, \Sigma_i') \tag{A.6}$$

□

Below we derive the EM algorithm update equations given in Section 2.3:

Solving for the intercept parameter $\mathbf{b}$

The value of $\mathbf{b}$ that solves $\frac{\partial Q}{\partial \mathbf{b}} = R^{-1} \sum_{i,j} (\mathbf{y}_{ij} - Z_{ij}\hat{\mathbf{u}}'_i - \mathbf{b}) = 0$ is given by

$$\begin{aligned} \mathbf{b} &= \frac{1}{\sum_i v_i} \sum_{i,j} (\mathbf{y}_{ij} - Z_{ij}\hat{\mathbf{u}}'_i) \\ &= \frac{1}{\sum_i v_i} \sum_{i,j} \left(\mathbf{y}_{ij} - \mathbf{a}s'_{ij}\right). \end{aligned} \tag{A.7}$$

Plugging in the expression for $\mathbf{a}$ from Equation A.9 yields

$$\mathbf{b} = \frac{\left(\sum_{i,j} \mathbf{y}_{ij}\right)\left(\sum_{i,j} \mathbf{q}_{ij}^T \Sigma'_i \mathbf{q}_{ij} + s'^2_{ij}\right) - \left(\sum_{i,j} \mathbf{y}_{ij} s'_{ij}\right)\left(\sum_{i,j} s'_{ij}\right)}{\left(\sum_i v_i\right)\left(\sum_{i,j} \mathbf{q}_{ij}^T \Sigma'_i \mathbf{q}_{ij} + s'^2_{ij}\right) - \left(\sum_{i,j} s'_{ij}\right)^2}. \tag{A.8}$$

Solving for the slope parameter $\mathbf{a}$

The value of $\mathbf{a}$ that solves $\frac{\partial Q}{\partial \mathbf{a}} = R^{-1}\left[\sum_{i,j}(\mathbf{y}_{ij} - \mathbf{b})s'_{ij} - \mathbf{a}\left(\sum_{i,j} \mathbf{q}_{ij}^T \Sigma'_i \mathbf{q}_{ij} + s'^2_{ij}\right)\right] = 0$ is given by

$$\mathbf{a} = \frac{1}{\sum_{i,j} \mathbf{q}_{ij}^T \Sigma'_i \mathbf{q}_{ij} + s'^2_{ij}} \sum_{i,j} (\mathbf{y}_{ij} - \mathbf{b}) s'_{ij}. \tag{A.9}$$

Plugging in the expression for $\mathbf{b}$ from Equation A.7 yields

$$\mathbf{a} = \frac{\left(\sum_i v_i\right)\left(\sum_{i,j} \mathbf{y}_{ij} s'_{ij}\right) - \left(\sum_{i,j} \mathbf{y}_{ij}\right)\left(\sum_{i,j} s'_{ij}\right)}{\left(\sum_i v_i\right)\left(\sum_{i,j} \mathbf{q}_{ij}^T \Sigma'_i \mathbf{q}_{ij} + s'^2_{ij}\right) - \left(\sum_{i,j} s'_{ij}\right)^2}. \tag{A.10}$$

Solving for the subject-specific variable mean parameter $\mathbf{m}$

The value of $\mathbf{m}$ that solves $\frac{\partial Q}{\partial \mathbf{m}} = -\sum_i V^{-1}\mathbf{m} + \sum_i V^{-1}\hat{\mathbf{u}}'_i = 0$ is given by

$$\mathbf{m} = \frac{1}{n} \sum_i \hat{\mathbf{u}}'_i. \tag{A.11}$$

Solving for the subject-specific variable variance parameter $\boldsymbol{\nu}$

We use numerical optimization to estimate $\boldsymbol{\nu}$:

$$\boldsymbol{\nu} = \arg\max_{\boldsymbol{\nu}} Q(\boldsymbol{\theta}, \boldsymbol{\theta}') \tag{A.12}$$

$$= \arg\min_{\boldsymbol{\nu}} \sum_i \left(\log|V| + (\hat{\mathbf{u}}'_i - \mathbf{m})^T V^{-1}(\hat{\mathbf{u}}'_i - \mathbf{m}) + \mathrm{Tr}\left(V^{-1}\Sigma'_i\right)\right) \tag{A.13}$$

We can reconstruct V by undoing the log-Cholesky parametrization steps using the estimated parameter vector $\boldsymbol{\nu} = [\nu_1, \nu_2, \nu_3]^T$ as

$$V = \begin{bmatrix} e^{\nu_1} & 0 \\ \nu_2 & e^{\nu_3} \end{bmatrix} \begin{bmatrix} e^{\nu_1} & \nu_2 \\ 0 & e^{\nu_3} \end{bmatrix}. \tag{A.14}$$

Solving for the noise covariance parameters $\boldsymbol{\lambda}$ and ρ

If $C = I_{K\times K}$ (i.e. $\rho = 0$ and fixed), then the diagonal elements λ_k of Λ can be estimated as:

$$\lambda_k = \sqrt{\frac{1}{\sum_i v_i}\sum_{i,j}\left[(y_{ijk} - a_k s'_{ij} - b_k)^2 + a_k^2 \mathbf{q}_{ij}^T \Sigma'_i \mathbf{q}_{ij}\right]}. \tag{A.15}$$

If C is not fixed to be the identity matrix, then in general it is not possible to obtain closed-form solutions for $\boldsymbol{\lambda}$ and ρ and we must use numerical optimization over Q. When $\boldsymbol{\lambda}$ is a high-dimensional vector, this is not feasible. Therefore, we simplify the parametrization of the noise covariance matrix R as $R(\lambda, \rho) = \lambda^2 \hat{\Lambda} C(\rho) \hat{\Lambda}$, where we now consider $\hat{\Lambda}$ as a fixed diagonal matrix with positive diagonal entries. We fix $\hat{\Lambda}$ at the estimate of Λ from the model with $C = I_{K\times K}$. ρ is the correlation matrix parameter as defined previously, and $\lambda > 0$ is a scaling parameter. Now we need to perform numerical optimization over Q to estimate only two parameters:

$$\lambda, \rho = \arg\max_{\lambda,\rho} Q(\boldsymbol{\theta}, \boldsymbol{\theta}'). \tag{A.16}$$

Note that since the EM algorithm does not require the maximization of $Q(\boldsymbol{\theta}, \boldsymbol{\theta}')$ but simply requires an increase in this function with each iteration, performing numerical optimization until convergence is not necessary.

Proposition A.2. *Consider the model given in Section 2.1.5 with $\boldsymbol{\theta} = \{\mathbf{m}, V, \mathbf{a}, \mathbf{b}, R\}$. The reparametrization $\boldsymbol{\theta}^* = \left\{ w\mathbf{m} + \begin{bmatrix} 0 \\ z \end{bmatrix}, w^2 V, \frac{1}{w}\mathbf{a}, \mathbf{b} - \frac{z}{w}\mathbf{a}, R \right\}$, where $w, z \in \mathbb{R}$, $w \neq 0$ yields an equivalent model.*

Proof. The incomplete log-likelihood is

$$\begin{aligned} \log f(\mathbf{y}; \boldsymbol{\theta}) = & \frac{1}{2}\sum_i \left(\log|2\pi\Sigma_i| - \log|2\pi V| - v_i \log|2\pi R|\right) \\ & -\frac{1}{2}\sum_{i,j}(\mathbf{y}_{ij} - \mathbf{b})^T R^{-1}(\mathbf{y}_{ij} - \mathbf{b}) \\ & -\frac{1}{2}\sum_i \mathbf{m}^T V^{-1}\mathbf{m} + \frac{1}{2}\sum_i \hat{\mathbf{u}}_i^T \Sigma_i^{-1} \hat{\mathbf{u}}_i. \end{aligned} \tag{A.17}$$

Here, we consider the case where $z = 0$ for algebraic simplicity. For the reparameterized model, we obtain $\hat{\mathbf{u}}_i^* = w\hat{\mathbf{u}}_i$ and $\Sigma_i^* = w^2\Sigma_i$. Therefore, each of the terms in $\log f(\mathbf{y}; \boldsymbol{\theta}^*)$ remain the same as those in $\log f(\mathbf{y}; \boldsymbol{\theta})$, yielding $\log f(\mathbf{y}; \boldsymbol{\theta}^*) = \log f(\mathbf{y}; \boldsymbol{\theta})$. □

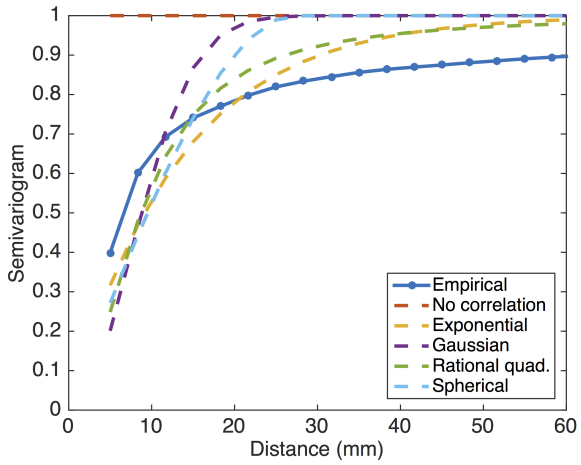

Inline Supplementary Figure 1: Preliminary analysis of the noise spatial correlation structure using the semivariogram. The empirical semivariogram computed using the residuals from the model where $C = I_{K \times K}$ is shown in blue. We fitted the semivariograms corresponding to the exponential, Gaussian, rational quadratic, and spherical correlation structures listed in Table 1 to the empirical semivariogram. The rational quadratic function had the best fit to the empirical curve based on the sum of squared error of the fitted semivariogram over a 100 mm distance calculated using 30 equidistant points.

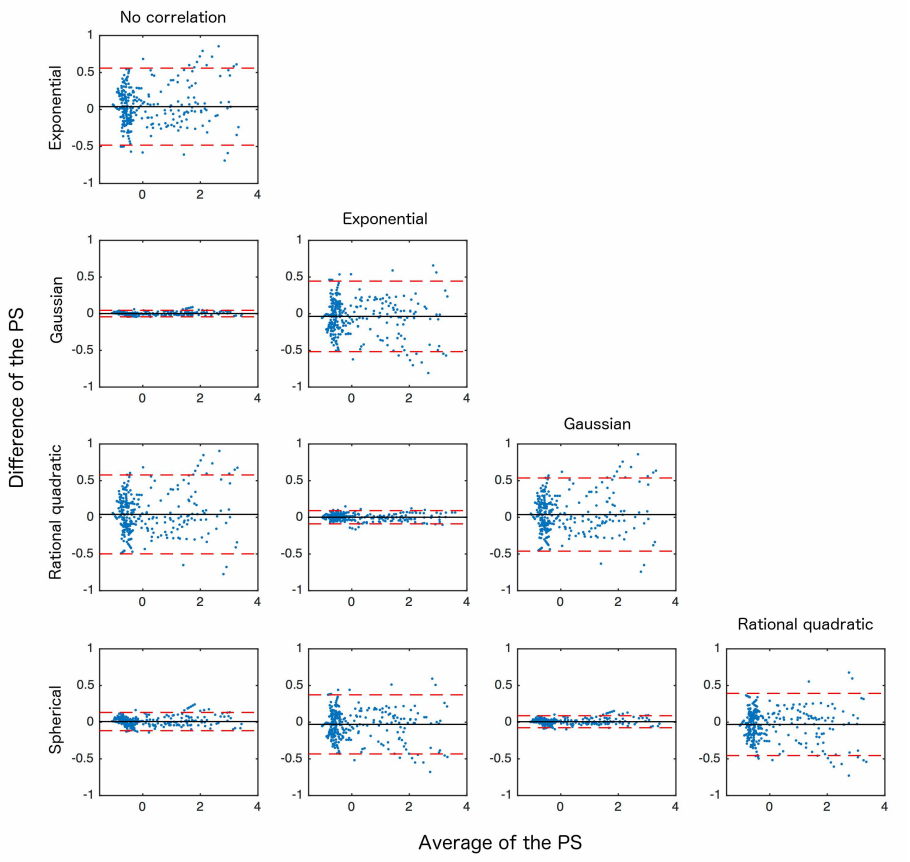

Inline Supplementary Figure 2: Bland-Altman plots comparing Aβ-PS values obtained using different spatial correlation models. We fitted the model using the "no correlation" structure ($C = I_{K\times K}$) and four spatial correlation structures (exponential, Gaussian, rational quadratic, and spherical). We then used Bland-Altman plots to assess the differences in PS computed using these correlation structures. In the Bland-Altman plots, the x-axis is the average of the PS values computed using the two methods being compared, and the y-axis is the difference. The labels above and to the left of the figures describe the spatial correlation structures being compared. Each blue dot corresponds to a visit. The solid black line indicates the mean difference between PS values computed using the methods, and the dashed red lines indicate its 95% confidence band. The Bland-Altman plots indicate that the "no correlation", Gaussian, and spherical structures yield similar PS values. Furthermore, the exponential and rational quadratic structures also yield similar PS values, but these differ from those obtained using the other three structures.

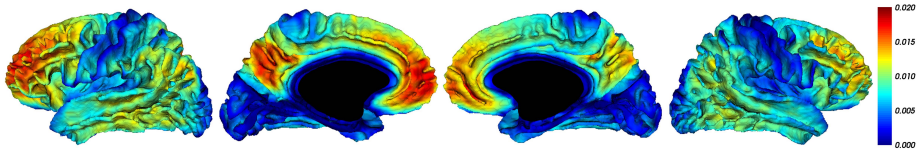

Inline Supplementary Figure 3: Slope parameters η_k (fixed effects) obtained from voxelwise LME model projected onto the cortical surface. DVR value at voxel k increases by η_k per year on average.

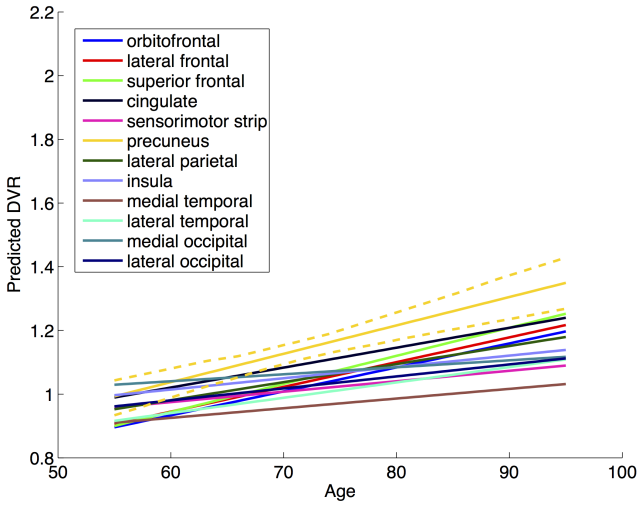

Inline Supplementary Figure 4: Regional trajectories obtained from the LME model as function of age. The LME model was used to make voxelwise predictions at a range of age values based on the estimated fixed effects, and these predictions were averaged within each ROI to obtain regional trajectories. The dashed lines indicate the 95% confidence bands for the cortical regions.

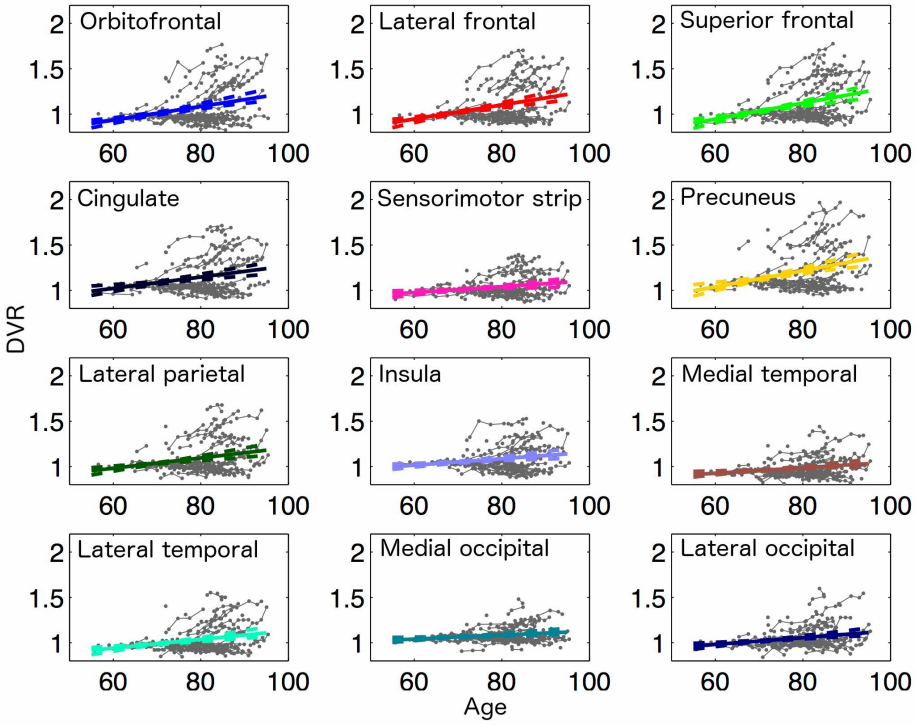

Inline Supplementary Figure 5: Regional trajectories obtained from the LME model as function of age. The LME model was used to make voxelwise predictions at a range of age values based on the estimated fixed effects, and these predictions were averaged within each ROI to obtain regional trajectories. The dashed lines indicate the 95% confidence bands for the cortical regions. Estimated trajectories with their 95% confidence bands are superimposed on observed longitudinal data (in gray).

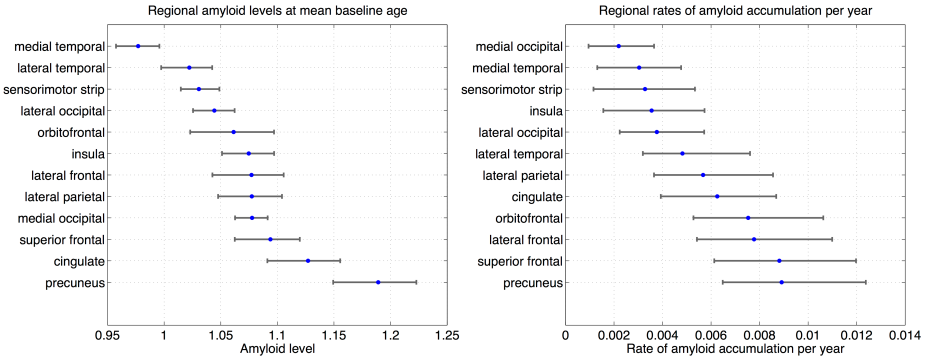

Inline Supplementary Figure 6: Comparison of levels of amyloid at age 77 (the mean baseline age of the sample) and rates of amyloid accumulation across cortical regions using results of the LME model. We used the fixed effect estimates to obtain voxelwise amyloid levels at age 77 as $\hat{y}_k = \eta_k \times 77 + \gamma_k$, and averaged $\hat{y}_k$ within each ROI to obtain regional amyloid levels at age 77. The fixed effects η_k were averaged within each ROI to obtain regional rates.

References

References

Avants, B.B., Epstein, C.L., Grossman, M., Gee, J.C., 2008. Symmetric diffeomorphic image registration with cross-correlation: evaluating automated labeling of elderly and neurodegenerative brain. Medical Image Analysis 12, 26–41.

Avants, B.B., Yushkevich, P., Pluta, J., Minkoff, D., Korczykowski, M., Detre, J., Gee, J.C., 2010. The optimal template effect in hippocampus studies of diseased populations. NeuroImage 49, 2457–2466.

Bateman, R.J., Xiong, C., Benzinger, T.L.S., Fagan, A.M., Goate, A., Fox, N.C., Marcus, D.S., Cairns, N.J., Xie, X., Blazey, T.M., Holtzman, D.M., Santacruz, A., Buckles, V., Oliver, A., Moulder, K., Aisen, P.S., Ghetti, B., Klunk, W.E., McDade, E., Martins, R.N., Masters, C.L., Mayeux, R., Ringman, J.M., Rossor, M.N., Schofield, P.R., Sperling, R.A., Salloway, S., Morris, J.C., 2012. Clinical and biomarker changes in dominantly inherited Alzheimer's disease. New England Journal of Medicine 367, 795–804.

Bernal-Rusiel, J.L., Greve, D.N., Reuter, M., Fischl, B., Sabuncu, M.R., 2012. Statistical analysis of longitudinal neuroimage data with linear mixed effects models. NeuroImage 66, 249–260.

Bernal-Rusiel, J.L., Reuter, M., Greve, D.N., Fischl, B., Sabuncu, M.R., 2013. Spatiotemporal linear mixed effects modeling for the mass-univariate analysis of longitudinal neuroimage data. NeuroImage 81, 358–370.

Bilgel, M., An, Y., Lang, A., Prince, J., Ferrucci, L., Jedynak, B., Resnick, S.M., 2014. Trajectories of Alzheimer disease-related cognitive measures in a longitudinal sample. Alzheimer's & Dementia 10, 735–742.

Bilgel, M., An, Y., Zhou, Y., Wong, D.F., Prince, J.L., Ferrucci, L., Resnick, S.M., 2015a. Individual estimates of age at detectable amyloid onset for risk factor assessment. Alzheimer's & Dementia [In press].

Bilgel, M., Carass, A., Resnick, S.M., Wong, D.F., Prince, J.L., 2015b. Deformation field correction for spatial normalization of PET images. NeuroImage 119, 152–163.

Bilgel, M., Jedynak, B., Wong, D.F., Resnick, S.M., Prince, J.L., 2015c. Temporal trajectory and progression score estimation from voxelwise longitudinal imaging measures: Application to amyloid imaging, in: Ourselin, S., Alexander, D.C., Westin, C.F., Cardoso, M.J. (Eds.), Lecture Notes in Computer Science (9123), Information Processing in Medical Imaging, pp. 424–436.

Biomarkers Definitions Working Group, 2001. Biomarkers and surrogate endpoints: Preferred definitions and conceptual framework. Clinical Pharmacology & Therapeutics 69, 89–95.

Braak, H., Braak, E., 1991. Neuropathological staging of Alzheimer-related changes. Acta Neuropathologica 82, 239–259.

Caroli, A., Frisoni, G.B., 2010. The dynamics of Alzheimer's disease biomarkers in the Alzheimer's Disease Neuroimaging Initiative cohort. Neurobiology of Aging 31, 1263–1274.

Cressie, N., Hawkins, D.M., 1980. Robust estimation of the variogram. Journal of the International Association of Mathematical Geology 12, 115–125.

Dale, A., Fischl, B., Sereno, M., 1999. Cortical surface-based analysis: I. Segmentation and surface reconstruction. NeuroImage 194, 179–194.

Delor, I., Charoin, J.E., Gieschke, R., Retout, S., Jacqmin, P., 2013. Modeling Alzheimer's disease progression using disease onset time and disease trajectory concepts applied to CDR-SOB scores from ADNI. CPT: Pharmacometrics & Systems Pharmacology 2, e78.

Desikan, R.S., Ségonne, F., Fischl, B., Quinn, B.T., Dickerson, B.C., Blacker, D., Buckner, R.L., Dale, A.M., Maguire, R.P., Hyman, B.T., Albert, M.S., Killiany, R.J., 2006. An automated labeling system for subdividing the human

cerebral cortex on MRI scans into gyral based regions of interest. NeuroImage 31, 968–980.

Donohue, M.C., Jacqmin-Gadda, H., Le Goff, M., Thomas, R.G., Raman, R., Gamst, A.C., Beckett, L.A., Jack, C.R., Weiner, M.W., Dartigues, J.F., Aisen, P.S., 2014. Estimating long-term multivariate progression from short-term data. Alzheimer's and Dementia 10, S400–S410.

Doody, R.S., Pavlik, V., Massman, P., Rountree, S., Darby, E., Chan, W., 2010. Predicting progression of Alzheimer's disease. Alzheimer's Research & Therapy 2.

Fonteijn, H.M., Modat, M., Clarkson, M.J., Barnes, J., Lehmann, M., Hobbs, N.Z., Scahill, R.I., Tabrizi, S.J., Ourselin, S., Fox, N.C., Alexander, D.C., 2012. An event-based model for disease progression and its application in familial Alzheimer's disease and Huntington's disease. NeuroImage 60, 1880–1889.

Galecki, A., Burzykowski, T., 2013. Linear model with fixed effects and correlated errors, in: Linear Mixed-Effects Models Using R. Springer New York, New York, NY. Springer Texts in Statistics. chapter 10, pp. 177–196.

Ito, K., Corrigan, B., Zhao, Q., French, J., Miller, R., Soares, H., Katz, E., Nicholas, T., Billing, B., Anziano, R., Fullerton, T., 2011. Disease progression model for cognitive deterioration from Alzheimer's Disease Neuroimaging Initiative database. Alzheimer's and Dementia 7, 151–160.

Jack, C.R., Knopman, D.S., Jagust, W.J., Petersen, R.C., Weiner, M.W., Aisen, P.S., Shaw, L.M., Vemuri, P., Wiste, H.J., Weigand, S.D., Lesnick, T.G., Pankratz, V.S., Donohue, M.C., Trojanowski, J.Q., 2013. Tracking pathophysiological processes in Alzheimer's disease: an updated hypothetical model of dynamic biomarkers. Lancet Neurology 12, 207–216.

Jack, C.R., Lowe, V.J., Senjem, M.L., Weigand, S.D., Kemp, B.J., Shiung, M.M., Knopman, D.S., Boeve, B.F., Klunk, W.E., Mathis, C.a., Petersen, R.C., 2008. ^{11}C PiB and structural MRI provide complementary information in imaging

of Alzheimer's disease and amnestic mild cognitive impairment. Brain 131, 665–680.

Jedynak, B.M., Lang, A., Liu, B., Katz, E., Zhang, Y., Wyman, B.T., Raunig, D., Jedynak, C.P., Caffo, B., Prince, J.L., 2012. A computational neurodegenerative disease progression score: Method and results with the Alzheimer's disease Neuroimaging Initiative cohort. NeuroImage 63, 1478–1486.

Jedynak, B.M., Liu, B., Lang, A., Gel, Y., Prince, J.L., 2014. A computational method for computing an Alzheimer's disease progression score; experiments and validation with the ADNI data set. Neurobiology of Aging 36 Supplement, S178–S184.

Jenkinson, M., Bannister, P., Brady, M., Smith, S., 2002. Improved optimization for the robust and accurate linear registration and motion correction of brain images. NeuroImage 17, 825–841.

Lindstrom, M., Bates, D., 1990. Nonlinear mixed effects models for repeated measures data. Biometrics 46, 673–687.

Mintun, M.A., Larossa, G.N., Sheline, Y.I., Dence, C.S., Lee, S.Y., Mach, R.H., Klunk, W.E., Mathis, C.A., DeKosky, S.T., Morris, J.C., 2006. [C]PIB in a nondemented population: potential antecedent marker of Alzheimer disease. Neurology 67, 446–452.

Pinheiro, J.C., Bates, D.M., 1996. Unconstrained parametrizations for variance-covariance matrices. Statistics and Computing 6, 289–296.

Resnick, S.M., Goldszal, A.F., Davatzikos, C., Golski, S., Kraut, M.A., Metter, E.J., Bryan, R.N., Zonderman, A.B., 2000. One-year age changes in MRI brain volumes in older adults. Cerebral Cortex 10, 464–472.

Rodrigue, K.M., Kennedy, K.M., Devous Sr., M.D., Rieck, J.R., Hebrank, A.C., Diaz-Arrastia, R., Mathews, D., Park, D.C., 2012. β-amyloid burden in healthy aging. Neurology 78, 387–395.

Schiratti, J.B., Allassonniere, S., Colliot, O., Durrleman, S., 2015a. Learning spatiotemporal trajectories from manifold-valued longitudinal data, in: Cortes, C., Lawrence, N., Lee, D., Sugiyama, M., Garnett, R. (Eds.), Advances in Neural Information Processing Systems 28, pp. 2395–2403.

Schiratti, J.B., Allassonniere, S., Routier, A., Colliot, O., Durrleman, S., 2015b. A mixed-effects model with time reparametrization for longitudinal univariate manifold-valued data, in: Ourselin, S., Alexander, D.C., Westin, C.F., Cardoso, M.J. (Eds.), Lecture Notes in Computer Science (9123), Information Processing in Medical Imaging, pp. 564–575.

Schmidt-Richberg, A., Guerrero, R., Ledig, C., Molina-Abril, H., Frangi, A.F., Rueckert, D., 2015. Multi-stage biomarker models for progression estimation in Alzheimer's disease. Lecture Notes in Computer Science (9123), Information Processing in Medical Imaging 9123, 387–398.

Schulam, P., Wigley, F., Saria, S., 2015. Clustering longitudinal clinical marker trajectories from electronic health data: Applications to phenotyping and endotype discovery, in: AAAI Conference on Artificial Intelligence.

Shock, N.W., Greulich, R.C., Andres, R., Arenberg, D., Costa Jr., P.T., Lakatta, E.G., Tobin, J.D., 1984. Normal human aging: The Baltimore Longitudinal Study of Aging. Technical Report. U.S. Government Printing Office. Washington, DC.

Sperling, R., Mormino, E., Johnson, K., 2014a. The evolution of preclinical Alzheimer's disease: implications for prevention trials. Neuron 84, 608–622.

Sperling, R.A., Rentz, D.M., Johnson, K.A., Karlawish, J., Donohue, M., Salmon, D.P., Aisen, P., 2014b. The A4 Study: Stopping AD Before Symptoms Begin? Science Translational Medicine 6, 228fs13.

Villain, N., Chételat, G., Grassiot, B., Bourgeat, P., Jones, G., Ellis, K.A., Ames, D., Martins, R.N., Eustache, F., Salvado, O., Masters, C.L., Rowe, C.C., Villemagne, V.L., 2012. Regional dynamics of amyloid-β deposition

in healthy elderly, mild cognitive impairment and Alzheimer's disease: a voxelwise PiB-PET longitudinal study. Brain 135, 2126–2139.

Villemagne, V.L., Burnham, S., Bourgeat, P., Brown, B., Ellis, K.A., Salvado, O., Szoeke, C., Macaulay, S.L., Martins, R., Maruff, P., Ames, D., Rowe, C.C., Masters, C.L., 2013. Amyloid β deposition, neurodegeneration, and cognitive decline in sporadic Alzheimer's disease: a prospective cohort study. Lancet Neurology 12, 357–367.

Villeneuve, S., Rabinovici, G.D., Cohn-Sheehy, B.I., Madison, C., Ayakta, N., Ghosh, P.M., La Joie, R., Arthur-Bentil, S.K., Vogel, J.W., Marks, S.M., Lehmann, M., Rosen, H.J., Reed, B., Olichney, J., Boxer, A.L., Miller, B.L., Borys, E., Jin, L.W., Huang, E.J., Grinberg, L.T., DeCarli, C., Seeley, W.W., Jagust, W., 2015. Existing Pittsburgh Compound-B positron emission tomography thresholds are too high: statistical and pathological evaluation. Brain 138, 2020–2033.

Yang, E., Farnum, M., Lobanov, V., Schultz, T., Raghavan, N., Samtani, M.N., Novak, G., Narayan, V., DiBernardo, A., 2011. Quantifying the pathophysiological timeline of Alzheimer's disease. Journal of Alzheimer's Disease 26, 745–753.

Younes, L., Albert, M., Miller, M.I., 2014. Inferring changepoint times of medial temporal lobe morphometric change in preclinical Alzheimer's disease. NeuroImage: Clinical 5, 178–187.

Young, A.L., Oxtoby, N.P., Daga, P., Cash, D.M., Fox, N.C., Ourselin, S., Schott, J.M., Alexander, D.C., 2014. A data-driven model of biomarker changes in sporadic Alzheimer's disease. Brain 137, 2564–2577.

Zhou, Y., Endres, C.J., Brašić, J.R., Huang, S.C., Wong, D.F., 2003. Linear regression with spatial constraint to generate parametric images of ligand-receptor dynamic PET studies with a simplified reference tissue model. NeuroImage 18, 975–989.

Ziegler, G., Penny, W.D., Ridgway, G.R., Ourselin, S., Friston, K.J., 2015. Estimating anatomical trajectories with Bayesian mixed-effects modeling. NeuroImage 121, 51–68.